arXiv:1512.08451v2 [cs.AI] 8 Jan 2016

GELATO and SAGE: An Integrated Framework for MS Annotation

Khalifeh AlJadda[1], Rene Ranzinger[2], Melody Porterfield[2], Mohammed Korayem[3], John A. Miller[1], Khaled M. Rasheed[1], Krys J. Kochut[1], and William S. York[2]

[1]Department of Computer Science, University of Georgia
[2]Complex Carbohydrate Research Center, University of Georgia
[3]School of Informatics and Computing, Indiana University

January 11, 2016

Abstract

Motivation: Several algorithms and tools have been developed to (semi) automate the process of glycan identification by interpreting Mass Spectrometric data. However, each has limitations when annotating MS^n data with thousands of MS spectra using uncurated public databases. Moreover, the existing tools are not designed to manage MS^n data where $n > 2$.

Results: Here, we propose a novel software package to automate the annotation of tandem MS data. This software consists of two major components. The first, is a free, semi-automated MS^n data interpreter called the Glycomic Elucidation and Annotation Tool (GELATO). This tool extends and automates the functionality of existing open source projects, namely, GlycoWorkbench (GWB) and GlycomeDB. The second is a machine learning model called Smart Anotation Enhancement Graph (SAGE), which learns the behavior of glycoanalysts to select annotations generated by GELATO that emulate human interpretation of the spectra.

Availability: GELATO is available within GRITS-Toolbox:
http://www.grits-toolbox.org/
SAGE to be available soon.

1 Introduction

Along with nucleic acids, proteins, and lipids, complex carbohydrates, also known as glycans, comprise the four major classes of macromolecules fundamental to all living systems [19]. Until recently, relatively little attention has been paid to studying glycans despite the major roles they play in diverse biological processes [8]. The emergence of the field of glycobiology is warranted

by the accumulated evidence for the role of glycans in cell growth and metastasis, cell-cell communication, and microbial pathogenesis. Almost all cells are coated with a dense layer of glycans and virtually all multicellular interactions take place in the context of this layer. Most proteins that are produced by eukaryotic cells for export or insertion into the cell membrane are glycosylated and proper glycosylation is often critical for their biological functions [3, 4]. Due to their complex structure, the potential information content encoded by glycans attached to proteins exceeds that of any other post-translational modification [14]. Comprehensive characterization of the glycans on glycoproteins has become an essential element for drug development, quality control, and basic biomedical research. However, glycan identification is much more difficult than protein identification, and de novo glycan sequencing is a proven NP-hard problem [16]. Glycans are more diverse than nucleic acids and proteins and peptides [10,19], mainly due to their branched structures. The linear structures of peptides and the availability of reliable peptide sequence databases facilitate their identification by tandem mass spectrometry (MS^n) [1] which generates a relatively complete series of high-intensity fragment ions with mass differences that correspond to specific amino acids and thus provides clear-cut information regarding the peptideâĂŹs amino acid sequence. In contrast, the branched nature of glycans often precludes the generation of readily interpretable ion series during the MS^n analysis . Furthermore, glycans are composed of monosaccharide building blocks that comprise isomeric sets (e.g., the set of all hexoses) whose members can be rarely distinguished by MS. Finally, annotation of MS data with structures from glycan databases is error-prone because existing glycan databases are incomplete, minimally curated, and frequently polluted with erroneous or irrelevant structures.

Current glycomics technology thus relies heavily on the manual interpretation of mass spectrometry datasets. As instrumentation improves, dataset size increases (e.g., 2000 or more mass spectra per biological sample), thus demanding more time for manual interpretation and reducing the number of samples that can be analyzed. This bottleneck is a major impediment blocking the application of glycoanalysis to a broad range of important biomedical investigations. Our evaluation of the existing tools for MS annotation revealed many limitations (Table 1). GlycoMod [7] is a free web-based tool intended for MS profile annotation but not for MS^n annotation. It annotates spectral features with compositions rather than structures. This tool supports limited types of chemical derivitisation and ionization adducts, with no support for neutral exchange. GlycoPeakfinder [12] is another web based tool that annotates both MS profiles and MS^n spectra. However, it only allows annotation of one spectrum at a time, and annotates with compositions rather than structures. GlycoWorkbench (GWB) [5] is freely-available software that annotates MS^n spectra using user defined structures or structures from a database. Its main limitation is that it can upload and process only one spectrum at a time, slowing down the data processing steps. Furthermore, incomplete chemical methylation of glycan structures [6] is not considered in the annotation process. The commercial tool SimGlycan® provides annotation of MS^n spectra with glycan structures from

an integrated database based on KEGG [11]. It facilitates high throughput analysis by uploading and annotating entire MS^n runs, but does not consider undermethylation of glycans. Here, we describe a new software package to facilitate the interpretation of MS^n data. This package consists of two components: GELATO, a freely available algorithm for the annotation of MS^n spectra of glycans, and SAGE, a machine learning model for refining GELATO annotations with trained expert knowledge.

2 Methods

Glycomic Elucidation and Annotation Tool (GELATO) is a freely available, semi-automated interpreter for MS^n of glycans designed and implemented at the Complex Carbohydrate Research Center (CCRC). GELATO extends the functionality of existing open source projects, namely, GlycoWorkbench (GWB) [5] and GlycomeDB [15].

The following extensions, which are not part of the GWB annotation algorithm, are implemented in GELATO:

1. Uploading the complete data set from an MS^n run at once rather than requiring each spectrum to be uploaded separately.

2. Specification of ionization adducts not predefined in our list on the basis of user-supplied information, including name, charge, and mass.

3. Annotation of multiply charged ions.

4. Annotation of ionic species generated by neutral ion exchange processes (e.g., replacement of H^+ with Na^+).

5. Annotation of ions generated by loss of small molecules (e.g., water or methanol).

6. Identification of ions arising from incompletely methylated glycans.

7. Ability to account for different fragmentation processes depending on the MS level or ion-fragmentation method.

8. Support for different accuracy (mass tolerance) settings for MS^1 and MS^n, consistent with various spectrometer setup protocols.

Confidence in the results produced by GELATO is increased by using a set of human curated databases generated from the glycan structure ontology, GlycO [17], as the source of glycan structures for the annotations of different types of glycans (e.g, N-glycans, O-glycans and Glycosphintolipids) [9]. Figure 1 and Figure 2 show the annotations produced by GELATO for MS^1 and MS^n, respectively. (The annotation workflow used by GELATO is provided in the supplementary document.) The GELATO annotation process starts by setting parameters that specify the types of ions (e.g., Na+ adducts), neutral ion

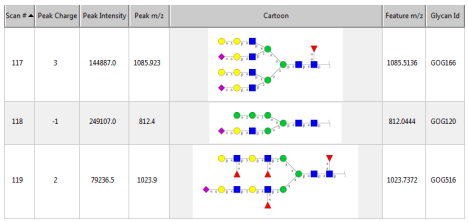

Figure 1: MS^1 annotations using GELATO

exchange events, and fragmentation processes (including, e.g., glycosidic bond cleavage and loss of water or methanol) that the user deems likely to generate the spectra being processed. Candidate glycan structures are then retrieved from the chosen database one at a time and each is checked to determine whether its calculated mass corresponds, within a specified tolerance, to the m/z of any precursor ion detected in MS^1. If a match is found, fragmentation of the candidate glycan is simulated in silico using the user-specified settings, and the calculated m/z values and structure of each theoretical fragment ion is saved.,The m/z values of the simulated fragment ions are compared with the m/z values of ions observed in the spectrum generated by fragmentation of the matching precursor ion, allowing observed (precursor and fragment) ions to be annotated with theoretical candidate structures. The resulting annotations are immediately serialized to a data file as they are assigned, making it possible for GELATO to run on a desktop computer or laptop with limited memory and CPU speed to handle MS data files containing hundreds of thousands of spectra. For MS^n spectra, this process is repeated recursively by choosing appropriate fragment ion annotations as candidate precursor structures for further *insilico* fragmentation and annotation of the spectra at the next MS level. GELATO generates and records all of the possible annotations based on these criteria for each processed spectrum without applying human expert knowledge or filtering the results. The annotations (each corresponding to a different candidate glycan structure) are then ranked using two complementary scoring metrics. The first, $score_c$, is calculated by dividing the number of fragment annotations for glycan G_x by the total number of ions in scan S_y, as shown in Equation 1 (Figure 3). Although the structurally relevant ions in a given spectrum usually exhibit high intensity, $score_c$ incorporates all annotated peaks regardless of their intensity. This can lead to less meaningful rankings if a large number of noise peaks are annotated because they happen to have m/z values that match the simulated fragments of the candidate glycan. This issue is addressed by calculating a second metric ($score_i$), which corresponds to the fraction of the total observed ion intensity in a fragment-ion spectrum that is derived from annotated peaks (Equation 2 in Figure 3). However, $score_i$ may also result in a misleading ranking if only one or few high intensity peaks are annotated. Therefore, both scores are provided to assist in user evaluation of the annotation results.

TThe GELATO annotation generation and ranking processes provide all the-

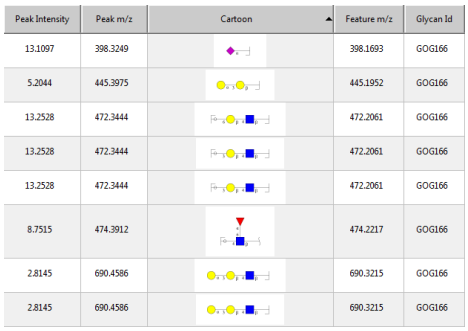

Peak Intensity	Peak m/z	Cartoon	Feature m/z	Glycan Id
13.1097	398.3249		398.1693	GOG166
5.2044	445.3975		445.1952	GOG166
13.2528	472.3444		472.2061	GOG166
13.2528	472.3444		472.2061	GOG166
13.2528	472.3444		472.2061	GOG166
8.7515	474.3912		474.2217	GOG166
2.8145	690.4586		690.3215	GOG166
2.8145	690.4586		690.3215	GOG166

Figure 2: MS^n annotations using GELATO

$$score_c(G_x|S_y) = \frac{num(AP_{xy})}{num(P_y)} \tag{1}$$

$$score_i(G_x|S_y) = \frac{\sum I(AP_{xy})}{\sum I(P_y)} \tag{2}$$

Figure 3: Scoring metrics calculated by GELATO. The variable num(AP_{xy}) is the number of peaks in spectrum S_y that are annotated by simulated fragments of G_x while $num(P_y)$ is the number of peaks in spectrum S_y. $I(AP_{xy})$ is the total intensity of peaks in spectrum S_y that are annotated by simulated fragments of G_x while $I(P_y)$ is the total intensity of peaks in spectrum S_y.

oretically possible annotations (for structures in the database), many of which may be meaningless. The user must therefore review the annotations and eliminate those that are judged to be incorrect, which can take considerable time and effort, especially when high throughput data are processed. In order to speed up the annotation of MS data, we thus implemented a machine learning model that identifies meaningful annotations based on human review patterns of previously processed datasets. This algorithm was implemented as a customization of the Probabilistic Graphical Model for Massive Hierarchical Data (PGMHD) [2]. The aim of the resulting tool, which we call Smart Annotation Enhancement Graph (SAGE), is to learn the annotation behavior of a user or group of users and apply this annotation behavior in subsequent annotations. The tool builds a probabilistic graph model that represents glycans and glycan fragments previously selected as meaningful annotations by the user and utilizes this graph to calculate the probability that the user would accept or reject a given annotation of the new dataset. SAGE can be used in two contexts, as an annotation tool, where it actually generates, scores and selects annotations, or as post-filtering tool, where it analyzes previously annotated spectra (e.g, processed by GELATO) to reject annotations that are unlikely to be accepted by the user. Training SAGE is straightforward and can be accomplished either in a single session or over many sessions. For example, spectra that have been annotated using GELATO might be reviewed by the user who selects a subset of the provided annotations that he/she judges to be correct. The selected annotations cann then be processed by SAGE to either build a new probabilistic graph model or integrate the new annotations into an existing model. The proposed learning algorithm for SAGE, shown in Algorithm 1, is designed to facilitate progressive learning, which has the following advantages:

1. Data required for training the model can be generated, evaluated by the user and processed in stages and at different times.

2. New training data is easily incorporated to extend the model without reprocessing data used in previous training sessions.

3. Training can be distributed over many sessions, eliminating the need for a single, prolonged session, which might fail and have to be repeated.

4. Recursive learning is possible, allowing the model itself to generate new training data by processing new MS data sets, provided that the new annotations are judged to be accurate by the user.

SAGE approaches MS annotation as a multi-label classification problem. Complete glycan structures that can annotate the spectra are the classes (into which observed spectra are assigned) while the fragments used to annotate the observed MS peaks are treated as features. Figure 4 illustrates the representation of MS data in SAGE. Each root node is labeled with a glycan structure and a specific m/z value, indicating that the user approved that structure during the training phase to annotate a precursor ion observed at that m/z. The nodes

```
Data: Annotated MS^n Spectra Using GELATO
Result: SAGE Instance
begin
    currentMSLevel = 0
    while currentMSLevel < maxMSLevel − 1 do
        foreach annotatedPrecursor ∈ currentMSLevel do
            if annotatedPrecursor.annotation exists in
            SAGE.currentLevelNodes then
                get sageNode where sageNode.annotation =
                annotatedPrecursor.annotation
                sageNode.frequency+ = 1
            else
                sageNode = newnode
                sageNode.frequency = 1
            end
            childrenLevel = currentMSLevel + 1
            foreach
            annotatedChildPeak ∈ annotatedPrecursor.peakList do
                foreach sageChildNode ∈ sageNode.children do
                    if annotatedChildPeak.annotation =
                    sageChildNode.annotation then
                        edge = edge(sagedNode, sageChildNode)
                        edge.frequency+ = 1
                    else
                        if childNode ∈ sage.childrenLevelNodes then
                            edge =
                            createNewEdge(sageNode, sageChildNode)
                            edge.frequency = 1
                        else
                            sageChildNode = newNode
                            sageChildNode.annotation =
                            annotatedChildPeak.annotation
                            sageChildNode.frequency = 1 edge =
                            createNewEdge(sageNode, sageChildNode)
                            edge.frequency = 1
                        end
                    end
                end
            end
        end
        currentMSLevel = currentMSLevel + 1
    end
end
```

Algorithm 1: Learning Algorithm for SAGE. $currentMSLevel$ represents the current MS level in the MS data we are processing, we start with level 0 which is related to MS level 1. $maxMSLevel$ is the highest level in the given MS data.

in the lower levels represent the fragments approved to annotate MS^2, MS^3, MS^n peaks. Edges are allowed between nodes at level i to nodes at level $i+1$ if and only if the ion observed and selected at level i decomposes to generate the fragment ion at level $i+1$. Numeric edge labels represent the number of times the parent node (the source of the edge) appears in the training data in association with the child node (the destination of the edge). After training, SAGE processes new (unannotated) data by associating the precursor ion of each MS scan with a set of root nodes (classes) and associating m/z values from the scan (peak list) with features in the trained model. The probabilistic classification score, $P(G_x|f_1, f_2, f_3, f_n)$ of each glycan (root node) is calculated given observation of those features in the scan. In order to optimize the search space, only those glycans forming quasi-molecular ions with m/z values that are within a specified tolerance from the precursor m/z for the given scan are considered. Figure 6 shows the annotation workflow for SAGE. The probabilistic score is calculated as described in [2]. For example, to use SAGE instance in Figure **??**, assume we have a spectra with the fragments F_1, F_3, F_7 and we would like to know which glycan structure in the root level is the best annotation. We can calculate the probabilistic score $P(G_1|F_1, F_3, F_7)$ and $P(G_2|F_1, F_3, F_7)$ as described in [2].

$$\begin{aligned} P(G_1|F_1, F_3, F_7) &= P(G_1|F_1, F_3) * P(F_3|F_7) \\ &= P(G_1|F_1) * P(G1|F_3) * (F_3|F_7) \\ &= 50/50 * 20/60 * 10/25 = 0.13 \end{aligned}$$

$$\begin{aligned} P(G_2|F_1, F_3, F_7) &= P(G_2|F_1, F_3) * P(F_3|F_7) \\ &= P(G_2|F_1) * P(G2|F_3) * (F_3|F_7) \\ &= 0.1 * 40/60 * 10/25 = 0.02 \end{aligned}$$

Since there is no edge between G_2 and F_1 even though it is a valid and possible fragment to G_2, to resolve the zero probability problem. Any fragment that cannot possibly be generated by fragmentation of a candidate glycan will not be included due to the selectivity of the model.

3 Discussion

The software described here combines the semi-automated MS annotation tool GELATO with the machine learning model SAGE to automate the process of MS annotation. As shown in Figure 5 A, the training of SAGE depends on the annotations generated by GELATO. The human analyst then examines and evaluates these annotations to select a subset that are deemed accurate and relevant. The selected annotations are then used to train SAGE. Figure 5 B, illustrates application of the trained SAGE for the de novo annotation of new spectra or improvement of the annotations previously generated by GELATO. This post-filtering process uses the knowledge learned by SAGE to eliminate the annotations that are most likely to be rejected by the user. SAGE thus

Table 1: Features of the existing MS annotation tools. Annotates MS^n, reflects the ability of the tools to annotate MS^n spectra, not just MS profile. Annotates with structures, means the tool uses database of glycans for annotation. Can find novel structures, means it can annotate using new glycan structures it never used before for annotation. Scores, define what type of scores the tools support. Statistical means scores calculated using statistical methods (summation, counting,etc.), while probabilistic means scores calculated based on probability distribution. Handle under-methylation, reflects the ability of the tools to consider under-methylation during the annotation process. Handle natural loss, reflects the ability to consider the natural loss in the annotation process. Average annotation time, reflects in general how long the annotation process require to be done. And availability, is either the tool is freely available or it is a commercial one.

Tool	Annotates MS^n	Annotates with structures	Can find novel structures	Scores	Handle Under-methylation	Handle Natural Loss	Uses human expert knowledge	Avg Annotation Time	Availability
GELATO	Yes	Yes	Yes	Statistical	Yes	Yes	No	Minutes	Free
SAGE	Yes	Yes	No	Probabilistic	No	No	Yes	Seconds	Free
SAGE and GELATO	Yes	Yes	Yes	Statistical and Probabilistic	Yes	Yes	Yes	Minutes	Free
GWB	Yes	Yes	Yes	Statistical	No	No	No	Minutes	Free
GlycoMod	No	Yes	Yes	Statistical	No	No	No	Minutes	Free
SimGlycan	Yes	Yes	Yes	Statistical	No	No	No	Hours	Commercial

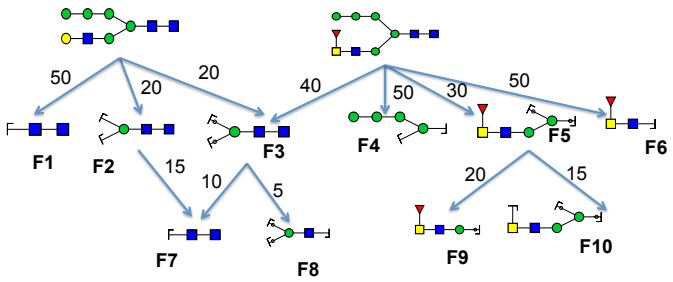

Figure 4: SAGE representing the MS data up to MS^3. The root nodes are glycans used to annotated precursors in the training data, while the nodes at lower levels represent fragments (F_i is the fragment Id) used to annotate the peaks in different MS levels. The edges represent the co-occurrence between the two nodes it connect while the number on the edge represent the frequency of that co-occurrence.

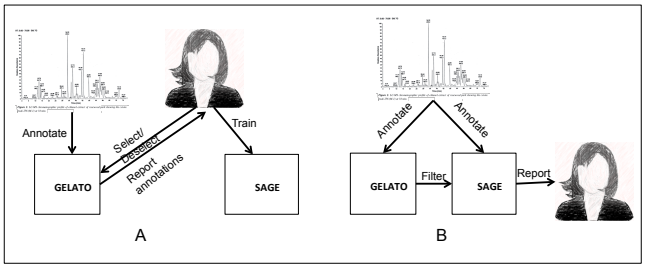

Figure 5: The integration between SAGE and GELATO. (A) GELATO annotates a given spectra, then a user select subset of those annotations and provide this final list of approved annotations to train SAGE. (B) the trained SAGE can be used either to annotate the given spectra or to filter out the annotations calculated by GELATO which most likely will not be selected by the user.

calculates a probabilistic score for each possible annotation (scan-glycan pairing), using previously learned knowledge of the annotation choices made by an analyst or group of analysts and reflecting the likelihood that the annotation would be accepted. The user can instruct SAGE to report the k top-scoring annotations. Otherwise, SAGE will report all the possible annotations ranked by their probabilistic score. This integrated framework addresses many of the shortcomings of existing tools, which are listed in Table 1.

As human knowledge is subjective, annotations that seem correct to one person may not be judged as correct by another person. Thus, SAGE not intended to provide a global knowledgebase that can be applied everywhere by everyone. Each instance of SAGE represents a particular analystâĂŹs knowledge, which he/she would like to apply to the annotation of new MS data sets. In this context, it is important to train new instances of SAGE to annotate MS data generated by analysis of different types of glycans (e.g., N-glycans vs. O-glycans) and the glycans from different taxonomic species, each of which are capable of generating their own unique collection of glycan structures.

4 Experiment and Results

To test our data-processing framework, we used GELATO to annotate 10 MS datasets generated during the glycoanalysis of stem cells. Annotations generated by GELATO were compared with annotations generated by SimGlycan, a commercial software tool specifically designed for the this purpose. For each spectrum tested, GELATO generated all the annotations generated by SimGlycan. However, GELATO was able to provide additional annotations that SimGlycan could not due to the following limitations. Firstly, SimGlycan cannot annotate scans whose precursor ions are not previously assigned a definite charge state (which is indicated in the input mzXML file). In contrast, GELATO identifies and annotates precursor ions with multiple charges by evaluating all

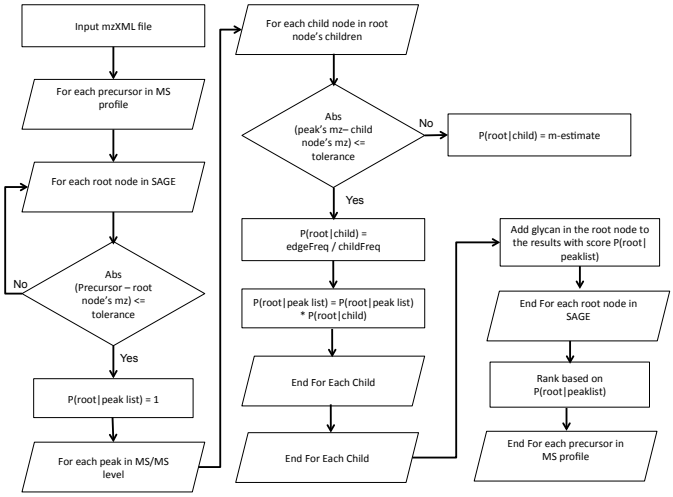
Input mzXML file
For each precursor in MS profile
For each root node in SAGE
Abs (Precursor – root node's mz) <= tolerance
No
Yes
P(root|peak list) = 1
For each peak in MS/MS level
For each child node in root node's children
Abs (peak's mz– child node's mz) <= tolerance
No
P(root|child) = m-estimate
Yes
P(root|child) = edgeFreq / childFreq
P(root|peak list) = P(root|peak list) * P(root|child)
End For Each Child
End For Each Child
Add glycan in the root node to the results with score P(root|peaklist)
End For each root node in SAGE
Rank based on P(root|peaklist)
End For each precursor in MS profile

Figure 6: The annotation workflow of SAGE. For more information about m-estimate, childFreq, edgeFreq and probabilistic based score please see [2]

charge states between 1 and the maximum charge (inclusive) specified by the user. Secondly, the database provided by and used by SimGlycan is incomplete. However, GELATO uses a custom glycan database generated from a highly curated source, such as the GlycO ontology. The integrated annotation framework was tested in a second experiment, where GELATO was used to annotate MS data obtained by glycoanalysis of 10 different samples from pancreatic cancer patient [13]. An MS expert reviewed the annotations generated by GELATO and selected a subset that she judged to be correct. Nine of the ten curated annotations were used to train SAGE, while the tenth was used for testing. The trained SAGE was used to generate de novo annotations of the tenth MS data set and these annotations were compared to those generated by GELATO and approved by the expert for the same dataset. This process was repeated 10 times with different test sets and training sets, and the average accuracy of the annotations was calculated. By these criteria, the annotations generated by SAGE were, on average, 98% accurate with 91% coverage (SAGE predicted 91% of the manually approved annotations of the test spectra). Since SAGE functions as a multi-label classifier and the MS annotation is multilabel classification problem, we compared SAGE with the top classifiers in machine learning. We used Mulan [18] which is an extension to the well-known machine learning library Weka to handle multi-label classification. Figure 7 shows that the SAGE implementation of PGMHD outperforms all the well-known classifiers in this challenging task. We are integrating our glycoanalysis annotation framework into the GRITS-Toolbox (http://www.grits-toolbox.org), a standalone application for the interpretation and annotation of glycomics MS data. Given that the GRITS-Toolbox is designed to run on a single workstation (desktop or lap-

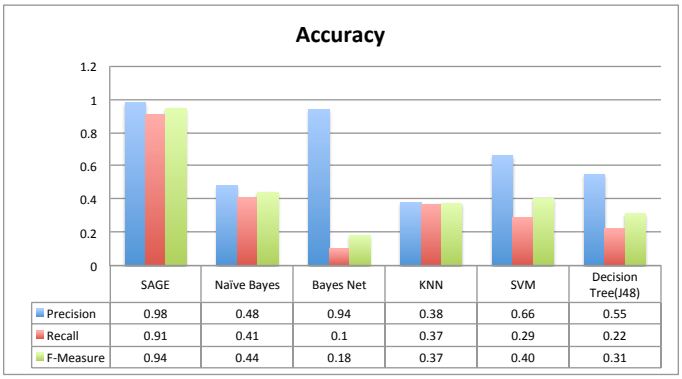

Figure 7: Precision, Recall, and F-Measure of the SAGE compared to the most popular classifiers

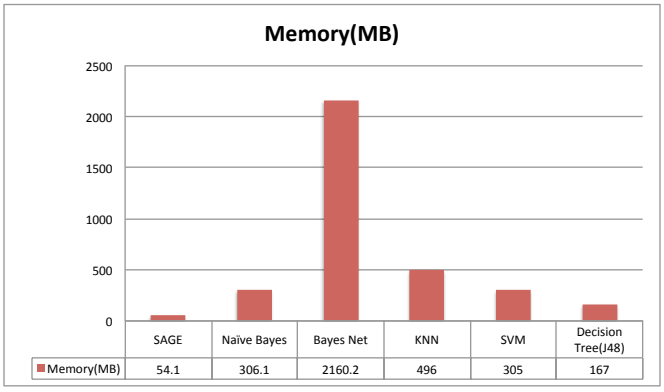

Figure 8: Main memory usage by SAGE compared to the other machine learning models for the training dataset

top) with limited memory, control of memory usage is a critical issue. In this context, we compared the memory usage of SAGE compared to the other classifiers. Figure 8 shows that the SAGE implementation of PGMHD used less memory (54 megabytes on average) than the other machine learning models to represent the training dataset. Another critical aspect of the training of and the annotation by SAGE is its time complexity. Figure 4.9 shows that, of all the tested machine learning models, SAGE required the least training time when using the same training dataset. Figure 10 compares the annotation time using different machine learning models,showing that SAGE is the third fastest in this respect.

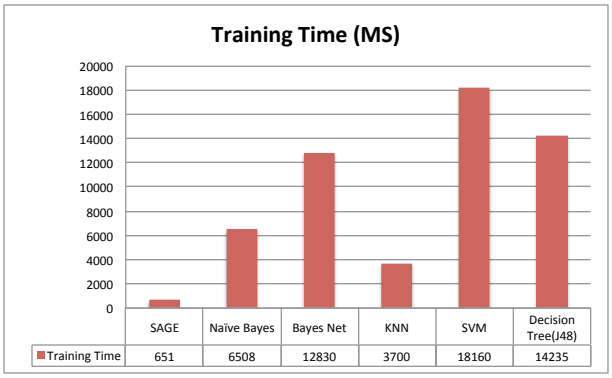

Figure 9: Training time for different models in millisecond. SAGE is the fastest model to learn and converge.

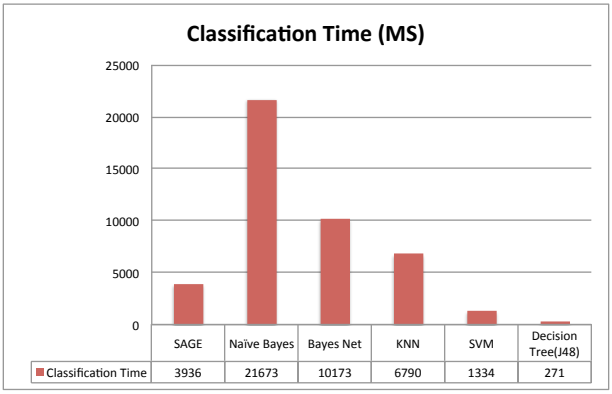

Figure 10: Annotation time for different models in millisecond. SAGE is the third fastest model.

5 Conclusion

We have designed and implemented an integrated software package for automated annotation MS data with glycan structures, This package consists of two major components, GELATO and SAGE. GELATO is a semi-automated annotation tool which is built upon two open source projects GlycoWorkbench and GlycomeDB. SAGE is a novel machine learning model that mimics the annotation patterns of an expert in MS interpretation to identify annotations of new data that are likely to be accepted or rejected by a human analyst . The current implementation of this framework utilizes GELATO to generate annotations that are used evaluated by a human. The selected annotations to train SAGE, which can then be used for either the de novo annotation of MS spectra or the improvement of annotations generated by GELATO. SAGE provides a probabilistic score that can be used as the basis for rejecting annotations that are likely to be incorrect.

Acknowledgement

We would like to thank Camilo Ortiz, and Trey Grainger from Careerbuilder.com for their contribution to the PGMHD model which we customize to SAGE. Also, we would like to thank Kiyoko Aoki Kinoshita from Soka University for the valuable time and deep discussion with her during the design of GELATO and SAGE. Many thanks to the team of EUROCarbDB for the well designed open source libraries which we built GELATO upon.

Funding This work was supported by the National Institute of General Medical Sciences, a part of the National Institutes of Health, funding the National Center for Biomedical Glycomics (8P41GM103490).

References

[1] Ruedi Aebersold and Matthias Mann. Mass spectrometry-based proteomics. *Nature*, 422(6928):198–207, 2003.

[2] Khalifeh AlJadda, Mohammed Korayem, Camilo Ortiz, Trey Grainger, John A Miller, and William S York. Pgmhd: A scalable probabilistic graphical model for massive hierarchical data problems. *arXiv preprint arXiv:1407.5656*, 2014.

[3] Rolf Apweiler, Henning Hermjakob, and Nathan Sharon. On the frequency of protein glycosylation, as deduced from analysis of the swiss-prot database. *Biochimica et Biophysica Acta (BBA)-General Subjects*, 1473(1):4–8, 1999.

[4] Susan A Brooks. Strategies for analysis of the glycosylation of proteins: Current status and future perspectives. *Molecular biotechnology*, 43(1):76–88, 2009.

[5] Alessio Ceroni, Kai Maass, Hildegard Geyer, Rudolf Geyer, Anne Dell, and Stuart M Haslam. Glycoworkbench: a tool for the computer-assisted annotation of mass spectra of glycansâĂă. *Journal of proteome research*, 7(4):1650–1659, 2008.

[6] Ping-Fu Cheng, Sergei Snovida, Ming-Yi Ho, Chu-Wen Cheng, Albert M Wu, and Kay-Hooi Khoo. Increasing the depth of mass spectrometry-based glycomic coverage by additional dimensions of sulfoglycomics and target analysis of permethylated glycans. *Analytical and bioanalytical chemistry*, 405(21):6683–6695, 2013.

[7] Catherine A Cooper, Elisabeth Gasteiger, and Nicolle H Packer. Glycomod–a software tool for determining glycosylation compositions from mass spectrometric data. *Proteomics*, 1(2):340–349, 2001.

[8] Raymond A Dwek. Glycobiology: toward understanding the function of sugars. *Chemical Reviews*, 96(2):683–720, 1996.

[9] Matthew Eavenson, Krys J Kochut, John A Miller, René Ranzinger, Michael Tiemeyer, Kazuhiro Aoki, and William S York. Qrator: A web-based curation tool for glycan structures. *Glycobiology*, page cwu090, 2014.

[10] David Goldberg, Marshall Bern, Bensheng Li, and Carlito B Lebrilla. Automatic determination of o-glycan structure from fragmentation spectra. *Journal of proteome research*, 5(6):1429–1434, 2006.

[11] Minoru Kanehisa. The kegg database. *Silico Simulation of Biological Processes*, 247:91–103, 2002.

[12] Kai Maass, René Ranzinger, Hildegard Geyer, Claus-Wilhelm von der Lieth, and Rudolf Geyer. ?glyco-peakfinder?–de novo composition analysis of glycoconjugates. *Proteomics*, 7(24):4435–4444, 2007.

[13] Mindy Porterfield, Peng Zhao, Haiyong Han, John Cunningham, Kazuhiro Aoki, Daniel D Von Hoff, Michael J Demeure, J Michael Pierce, Michael Tiemeyer, and Lance Wells. Discrimination between adenocarcinoma and normal pancreatic ductal fluid by proteomic and glycomic analysis. *Journal of proteome research*, 13(2):395–407, 2013.

[14] Christoph Rademacher and James C Paulson. Glycan fingerprints: calculating diversity in glycan libraries. *ACS chemical biology*, 7(5):829–834, 2012.

[15] René Ranzinger, Stephan Herget, Claus-Wilhelm von der Lieth, and Martin Frank. Glycomedb?a unified database for carbohydrate structures. *Nucleic acids research*, 39(suppl 1):D373–D376, 2011.

[16] Haixu Tang, Yehia Mechref, and Milos V Novotny. Automated interpretation of ms/ms spectra of oligosaccharides. *Bioinformatics*, 21(suppl 1):i431–i439, 2005.

[17] Christopher J Thomas, Amit P Sheth, and William S York. Modular ontology design using canonical building blocks in the biochemistry domain. *Frontiers in Artificial Intelligence and Applications*, 150:115, 2006.

[18] Grigorios Tsoumakas, Eleftherios Spyromitros-Xioufis, Jozef Vilcek, and Ioannis Vlahavas. Mulan: A java library for multi-label learning. *The Journal of Machine Learning Research*, 12:2411–2414, 2011.

[19] Daniel B Werz, René Ranzinger, Stephan Herget, Alexander Adibekian, Claus-Wilhelm von der Lieth, and Peter H Seeberger. Exploring the structural diversity of mammalian carbohydrates (?glycospace?) by statistical databank analysis. *ACS chemical biology*, 2(10):685–691, 2007.